I. INTRODUCTION

The electronics designed for the D0 Muon Upgrade [1,2] includes three detector subsystems and two trigger systems. This document describes front-end and readout electronics. The muon trigger systems are described in the separate documents [3,4]. The front ends and readout hardware are synchronized by means of timing signals broadcast from the D0 Trigger Framework. The front-end electronics have continuously running digitizers and two levels of buffering resulting in nearly deadtimeless operation. The raw data is corrected and formatted by 16-bit fixed point DSP processors. These processors also perform control of the data buffering. The data transfer from the front-end electronics located on the detector platform is performed by serial links running at 160 Mbit/s. The electronics for the muon system includes 9,500 channels for proportional drift tubes (PDT), 6,000 channels for various types of scintillation counters, and 50,000 channels for mini-drift tubes [5,6]. A major design goal was to develop a unified strategy for readout for all the detector subsystems.

All subsystems that can form a trigger send trigger data to the Trigger Framework (TFW) which is the source of all trigger decisions. These decisions are made on two levels based on inputs from the Level 1 (L1) and Level 2 (L2) trigger systems. The L1 trigger systems generate trigger information synchronously with the beam crossings while the L2 systems operate asynchronously and have an indeterminate decision time within some limit (A third level of trigger is a software filter using complete event information whose decisions are independent of the TFW). The Fermilab accelerator, running in the collider mode, uses the Tevatron RF frequency of 53.104 MHz to synchronize colliding beams. A clock frequency derived from the Tevatron RF divided by seven is used to synchronize the beam crossings within the accelerator. This frequency is the basic frequency for transferring L1 trigger information to and trigger decisions from the D0 Trigger Framework.

All the muon subsystem front ends use a strategy of continuously running digitizers for time and charge measurements. The digitizer outputs are connected to digital delays providing the necessary delay for the event data before a trigger decision is returned. Upon arrival from the TFW of an L1 accept decision, data is transferred to the first level of buffering. Both digitizer and delays run synchronously with the Tevatron RF. The front ends are globally synchronized to the beam crossings and each other by the TFW. It distributes timing and control signals including a copy of the Tevatron RF, L1 and L2 decisions and trigger numbers to the front-end electronics near the detector via the Geographic Sectors (GS) located in the movable counting house. These numbers are used to identify the beam crossing corresponding to the trigger decision and must be consistent across the detector.

A. Muon Detector Subsystems

The PDTs (9,500 channels total) have up to 96 wires per chamber. Attached to these chambers are three or four 24-channel Front-End Boards (FEB) and one Control Board (CB). The FEBs digitize the time of arrival of the wire signal and the charge of the pad electrodes. The digital delays and L1 buffers are also located on this board. Each FEB is connected to the CB via an 18 bit uni-directional data bus. The CB has a readout controller

which fetches data sequentially from each L1 FIFO upon receiving an L1 accept and stores it in DSP memory.

The scintillation counter subsystem (6,000 channels total) includes 48-channel 9U VME Scintillator Front-End cards (SFE). These cards measure arrival time of the scintillation counter signals. There are three charge integrators and ADCs each of which can be attached to a particular channel for use as a photomultiplier gain monitor. Since the gain of a particular channel is not checked on an event by event basis, one ADC can serve multiple channels by means of analog multiplexing. The same time digitizer chip [7] is used by both the FEBs and the SFEs. The L1 buffers are read out by a DSP based readout controller on a custom high speed data bus using the J3 backplane. Each VME crate houses up to ten SFE cards, the readout controller card, and a Motorola 680xx based VME processor to provide parameter downloading and testing.

The Mini-Drift Tubes (MDT, 50,000 channels total) have their electronics located in 192-channel 9U VME cards. The Mini-drift tube Digitizing Cards (MDC) perform a low resolution (18.8 ns/bin) measurement of the drift time, improving the coordinate resolution of the tubes to a few millimeters compared to the 10 mm tube diameter. The MDC also has a digital pipeline and L1 buffers to store event data. They also are read out by a specialized DSP based readout controller via the J3 backplane. The MDT crates can accommodate up to twelve MDCs, a readout controller and a VME processor to perform tasks similar to those of the scintillator system.

B. Trigger Architecture

The muon system must provide both L1 and L2 trigger information to the TFW (Fig. I-1) as opposed to some other D0 detectors which send trigger information only to L2. The muon L1 trigger system receives data from the various front-end systems and provides information every 132 ns to the TFW. Due to the distributed nature of the muon system there are a large number of L1 trigger inputs (over 150). The muon L1 concentrator cards perform data compression and an initial stage of pattern recognition to reduce the data flow into the L1 trigger system to a reasonable level. The TFW analyzes trigger information and generates an L1 decision which is distributed via high speed serial links to the Geographic Sectors. The TFW attaches a unique trigger number to each L1 decision for event synchronization. This is discussed in more detail in the next section. Each muon front end also provides L2 information to the Level 2 trigger system. Upon receiving an L1 accept, the front ends not only buffer the event, they also send a portion of the event to the L2 trigger system via serial links.

A GS consists of a VME crate containing the D0 standard interface to the DAQ called the VME Buffer Driver (VBD) and a control and timing interface between the TFW and the front ends connected to it. The front-end electronics and L1 trigger system are located on the D0 detector platform and connected to the readout crates by both twist and flat and ribbon coaxial cables. The Muon Fanout Card (MFC) distributes timing, control and trigger information by means of custom J2 connections on the crate backplane. It also receives status and error signals from the Muon Readout Cards (MRC) located in the crate and transfers this information back to the TFW.

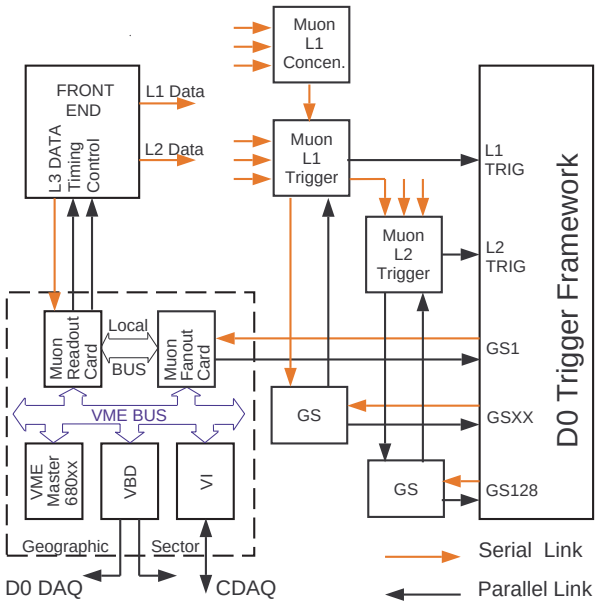

Fig. I-1. D0 muon trigger architecture. GS - geographic sector.

The Muon Readout Card is an intermediate stage between the MFC and two front ends. It sends timing, control and trigger information, and receives status and error signals from the front ends. It also has internal memory for buffering one event. The MRC and MFC are the standard muon GS interface cards. The data transfers between the VBD and MRCs are initialized by the 680xx processor, which also performs trigger number checking and event header formatting. In local mode, it is used for diagnostics and local data taking and processing. The Vertical Interconnect (VI) is a part of the D0 slow control and monitoring system. It affords remote access to the crate VME address space.

C. Event Synchronization

A major concern for buffered data acquisition systems is synchronization of the data comprising the events. We propose to use several levels of checking to identify errors and simplify debugging. The structure proposed for colliding beams in the Tevatron consists of three superbunches separated by abort gaps. The beam crossing intervals may be 132 or 396 ns, but in either case, the TFW synchronization time unit will be 132 ns. Each front-end receives a copy of the Tevatron RF and a special reset signal (First Crossing, FC, Fig. I-2) synchronized to it and arriving once per beam turn at a fixed offset from the beginning of the first collision in the superbunch. Due to the difference in the length of the cables connecting the front ends to the MCH, each front-end controller has the ability to adjust the timing of these signals over a range of approximately 250 ns with at least 18 ns of resolution. The clock and reset signals are used by the front ends to run two counters which together provide a unique identifier for each collision. One eight bit counter is clocked at 53/7 MHz and generates a crossing number. This counter is preset to one at the start of each turn. A second counter, 16 bits wide, clocked by the FC signal, generates a turn number. These counter values are delayed and read out in the same manner as data.

These two numbers are appended to events and are used by the data acquisition system for checks of event synchronization. In the muon system the crossing number simplifies the task of timing adjustment thus increasing system reliability. If there is no match between the crossing number emerging from the pipeline and the trigger number sent by the TFW, an error is generated.

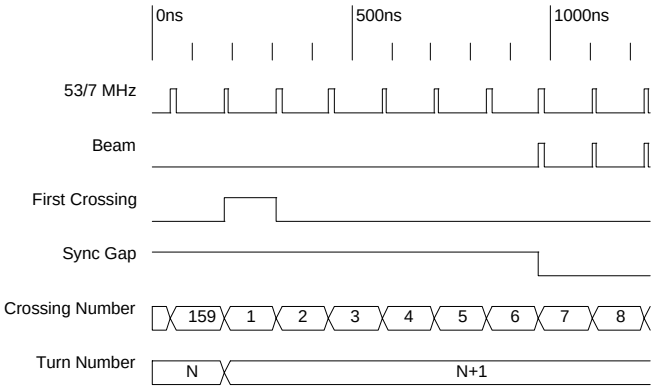

Fig. I-2. Timing diagram showing relevant signals at first crossing.

A global initialization signal (INIT) is generated by the TFW to synchronize all sub-systems before data taking begins. An INIT Acknowledge (INACK) is returned by the GSs when they finish their local initialization. When all the GSs return INACK, the TFW releases INIT and begins issuing trigger decisions. This method of synchronization allows for rapid re-synchronization and quick detection of any sub-systems that fail to initialize properly.

D. System Timing Synchronization.

Orbit timing signals are required by all of the muon front ends in order to keep a count of beam crossings and the number of beam orbits. These signals must be timed with a accuracy of one 18.8 ns RF clock interval. For those muon subsystems employing the TMC chip, an unadulterated copy of 53MHz is transmitted from the MCH to the front ends. There are a limited number of high bandwidth cables between the MCH and the front ends. Trigger decisions are sent to the muon readout unit by means of a 25 pair ribbon cable running from the MRC in the MCH to the collision hall. Peaking networks are fitted to each signal pair to increase the effective bandwidth of the cable, since the intrinsic bandwidth of this cable is less than 7MHz. To keep EMI radiation to a minimum, current mode transmission is used. Encoded timing information is sent on one section of a four conductor ribbon coaxial cable also running from the MRC to the platform. The other conductors are used for L2 trigger data, L3 data, and a 53MHz clock. To enforce a balance between center conductor and shield currents, and thus reduce EMI, transformer coupling is used. This has the additional feature of breaking any ground loops between the MCH and the platform.

As it happens, commercial ICs that have become available recently make this task straightforward. The encoding scheme must be fast enough to separately define a rising and falling edge within 132 ns.

The simplest method to send a data stream is to send a synchronous serial data stream on one line and the corresponding clock on another. Two problems arise from this method. First, the bandwidth of the coaxial cable running from the MCH to the front ends is small compared to 18.8 ns. Cable compensation is therefore required. This can be accomplished with a high speed amplifier and a peaking network matched to the cable. However an active cable equalizer chip available from National Semiconductor does this job with only a single external component required to specify the time constant of the cable. This device has the requirement that there be at least one signal transition every 50 ns. For this reason, the data must be encoded into a clock stream. The second problem is the very small setup and hold requirements for 18.8 ns timing resolution. The relative phase of the clock and data will change after being transmitted through separate drivers, cables and receivers. By sending clock and data together, all the channel delays apply equally to both clock and data and their relative phase is invariant.

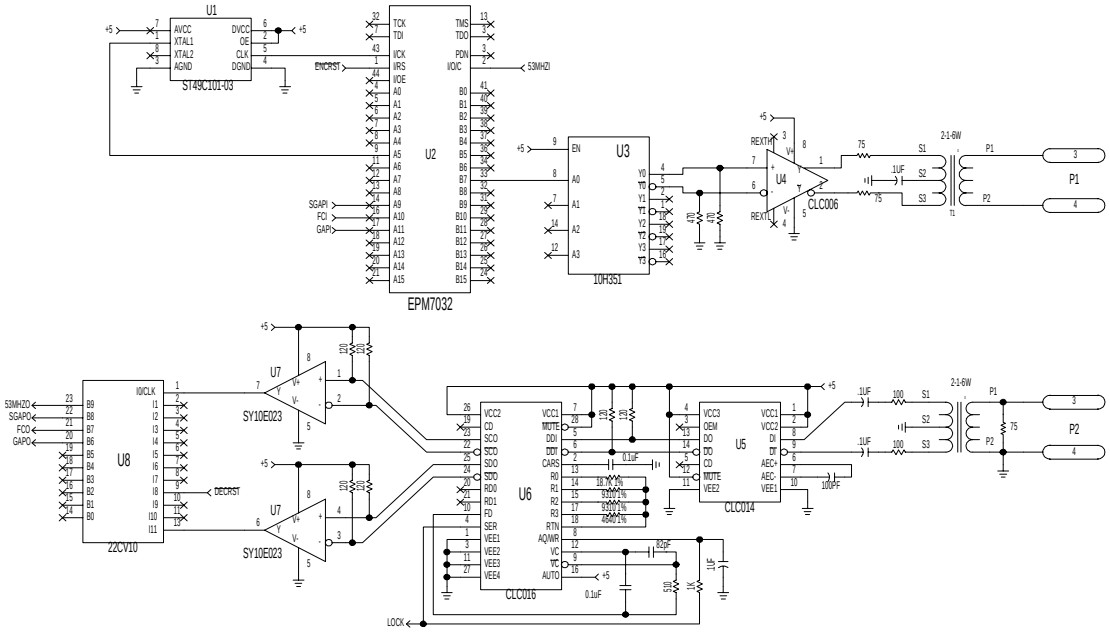

Fig. I-3. Encoder and decoder schematic diagram.

The preferred method for generating data sequences is the use of a synchronous state machine. A logical choice for the clock frequency to be sent with the data is 53 MHz, the basic frequency for the front ends. For a state machine to generate 53 MHz, it must be clocked at 106 MHz. A commercial CMOS TTL clock multiplier chip is used for this purpose. The Startech ST49C101 is available pre-programmed to a number of different multiplier values. In the x8 configuration, the specified maximum frequency of operation is 130 MHz. In order to use this device, the 53 MHz clock stream arriving from the Trigger Framework is first divided by 4 then multiplied by 8. The phase of the output encoded clock stream must be fixed with respect to the incoming clock, so the state machine generating the encoded clock stream must have the original 53 MHz signal as one of its inputs.

The TTL data is converted to PECL then sent through a cable driver. The cable equalizer (National CLC014) is attached at the receive end, which produces a PECL output. A commercial data separator chip (National CLC016) produces a data and clock stream. The PECL levels are converted to TTL and connected to a synchronous state machine decoder section. A schematic of the encoder and decoder section is shown in Fig. I-3. Because AC coupling is used, the encoded patterns must be DC balanced. The cable equalizer has better noise margins with DC balanced data as well. If there are three signals

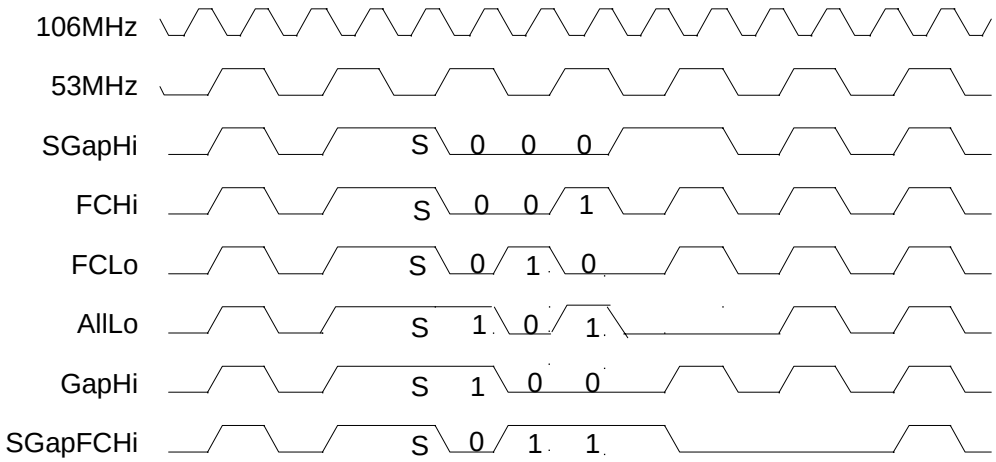

Fig. I-4. Encoder patterns.

each with a rising and falling edge, six unique patterns would be required to define these edges. We know, however, that the trailing edges of Gap and Sync Gap will never be coincident, so we can use the same pattern for the trailing edge of these two signals. There is a possibility that one edge of First Crossing may be coincident with one edge of Sync Gap. We can use the pattern that forces Sync Gap and Gap down to force First Crossing down as well. This definition will allow for the trailing edge of First Crossing to be coincident with the trailing edge of Sync Gap. That leaves then the case where the rising edge of Sync Gap and First Crossing are coincident. Fig. I-4 summarizes the waveforms for the different patterns.

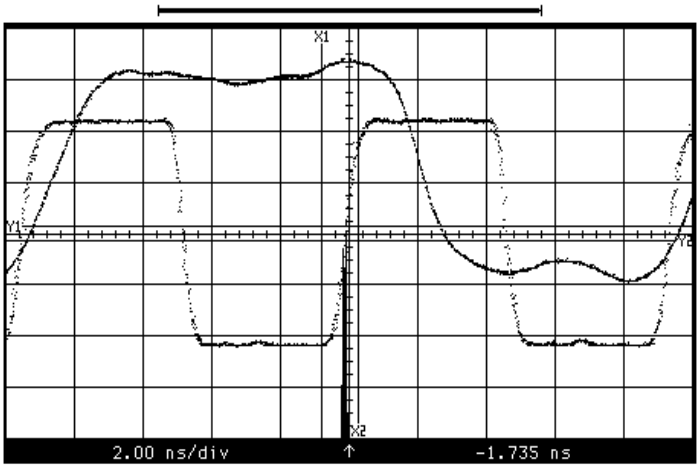

Fig. I-5. Input clock and recovered clock.

The logic for the two sections is straightforward. The encoder section is somewhat more complicated due to the requirement that a DC balanced sequence be set. Both encoder and decoder are essentially shift registers, where the encoder shifts a three bit pattern out and the decoder shifts a three bit pattern in. The patterns are actually eight bits long due to the balancing. An example files are written in VHDL and Altera AHDL along with some simulation waveforms. It should be possible to adapt these equations to a particular logic design language. An encoder and decoder test board was built using circuitry shown in Fig. I-3. Traces of the 53MHz input clock fed to the encoder and the 106MHz recovered PECL clock output of the data separator are shown in Fig. I-5. Fig. I-6 shows a jitter measured for the recovered clock.

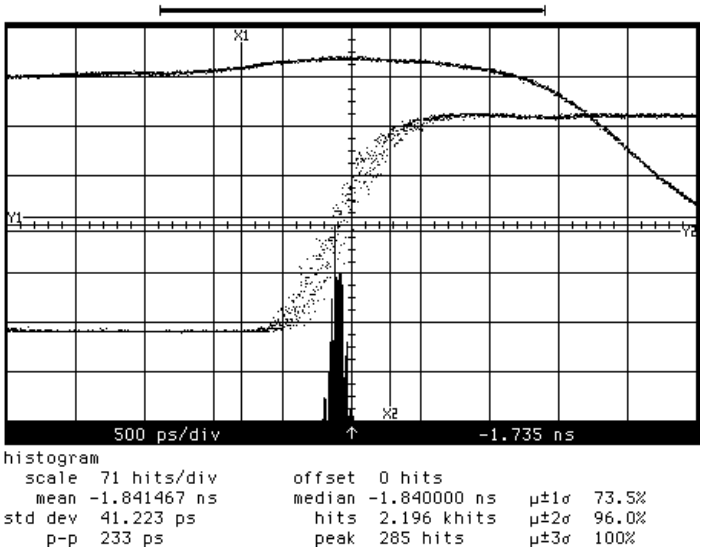

Fig. I-6. Jitter of recovered clock.

E. Level 1 Serial Data Transmitter

(a) Introduction.

All muon front ends use an unified serial transmitter daughter board to transfer trigger bits to the Level 1 trigger system. The Serial Transmitter Daughter Board accepts 16 bit parallel data at rates from 47.8 MHz to 58.4 MHz and creates a serial data stream at 20 times the parallel word rate that is 8b/10b encoded.

(b) Parallel Inputs.

The parallel data is input through a 30 pin connector (Samtec P/N TFM-115-02-S-D-LC), all inputs are compatible with both 3.3v and 5v logic inputs. The inputs are shown in Table I-1, Word_Clock is a free-running clock at the RF frequency and all other inputs, including Enable, are referenced to this clock. The input levels are compatible with either 5V or 3.3V logic families, depending on how power is applied to the connector. Pin 27 can be connected to either 3.3V or 5V to set the input voltage levels.

All inputs have a setup time (t_{su}) of 8ns with respect to Word_Clock and a hold time (t_h) of 2ns with respect to Word_Clock. The board transmits the SYNC character (K28.5) when the Enable signal is low, encoded data when the Enable Signal is high and odd

longitudinal parity when both the Enable Signal and Parity_Enable signals are high (see Appendix A). Parity_Enable causes odd longitudinal parity, calculated starting from the last Parity_Enable or Enable signal, to be transmitted in place of data. The parity word is typically added as the seventh word after the six data words sent for a bunch crossing, but can be added over any length of words at the discretion of the user. Serial Data will start appearing at the serial output 4 clock cycles after it is input to the board.

Table I-1. Parallel inputs.

Pin #	Label	Description
1	GND	Ground
2	GND	Ground
3	Data_In_0	LSB of Input Data
4	Data_In_1	Input Data
5	Data_In_2	“
6	Data_In_3	“
7	Data_In_4	“
8	Data_In_5	“
9	Data_In_6	“
10	Data_In_7	“
11	Data_In_8	“
12	Data_In_9	“
13	Data_In_10	“
14	Data_In_11	“
15	Data_In _12	“
16	Data_In_13	“
17	Data_In_14	“
18	Data_In_15	MSB of Input Data
19	Fast_OR	Fast Or Output
20	Word_Clock	RF Clock
21	Enable	High = Data Enabled
22	Parity_Enable	High = Transmit Parity
23	Spare	
24	GND	Ground
25	+5V	+5 volt Power Supply
26	+5V	+5 volt Power Supply
27	+3.3V or +5v	Output Power Supply
28	GND	Ground
29	+3,3V	+3.3 volt Power Supply
30	+3.3V	+3.3 volt Power Supply

(c) Serial Outputs

The Serial Data is output through a MCX 50Ω connector mounted directly on the board.

(d) JTAG/Programming Connector

A 10 pin connector (Samtec P/N TFM-105-02-S-D-A) is provided for those users who wish to use the JTAG Boundary Scan Test capabilities of the Altera EPLD that drives the output lines of the Serial Transmitter Daughter Board. Users that do not wish to

reprogram the Serial Transmitter card in system can ignore this connector. Table I-2 shows the pin assignments for this connector.

Table I-2. JTAG connections.

Pin #	Label	Description
1	TCK	Test Clock In
2	GND	Ground
3	TDO	Test Data Out
4	Vcc	+5v
5	TMS	Test Mode Select
6	NC	
7	NC	
8	NC	
9	TDI	Test Data In
10	GND	Ground

(e) Timing Requirements

Fig. I-7 shows the relationship between the Input Data, Parity_Enable and the clock. Words 1 to 6 contain the data for the bunch crossing and word 7 will contain the parity word that is generated on the Serial Transmitter Daughter Board. Note that the Serial Transmitter Daughter Board will generate odd parity for all the words since the last Parity_Enable signal, so it is possible to generate parity of messages of any length.

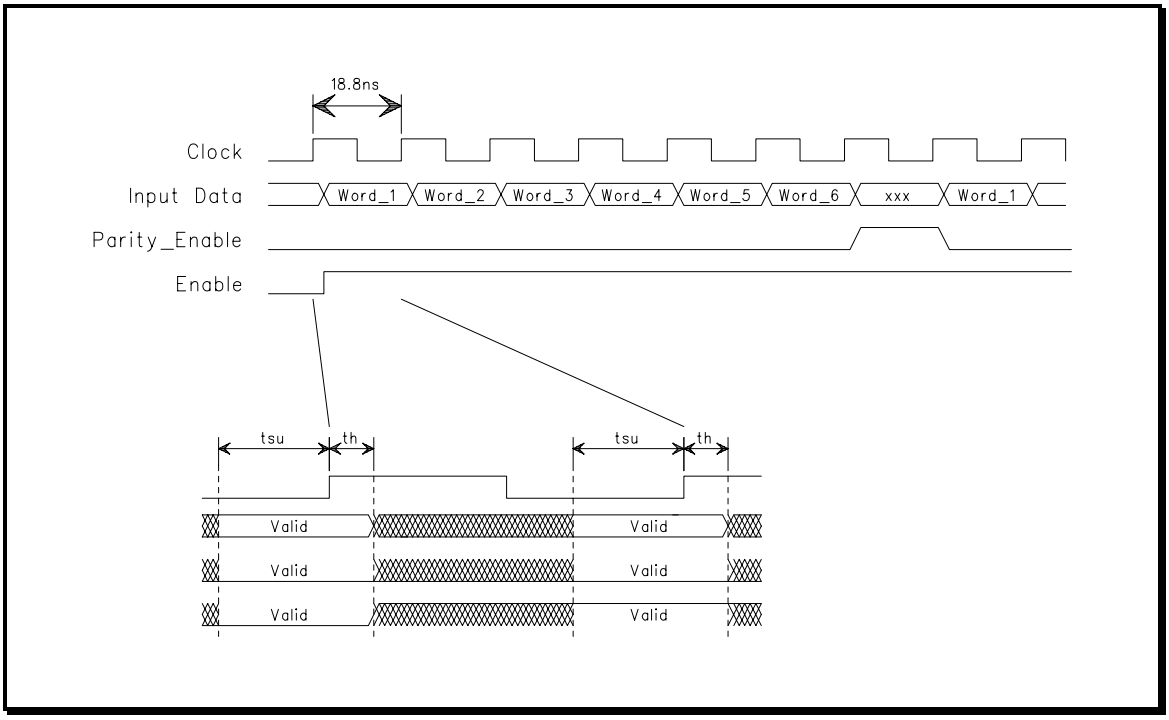

Fig. I-7. Input timing.

(f) Mechanical Layout

The Serial Transmitter Daughter Board is 1.5” x 2.2” and has 4 mounting holes as well as the connectors that attach to the daughter board. Fig. I-8 shows the mechanical arrangement of the board.

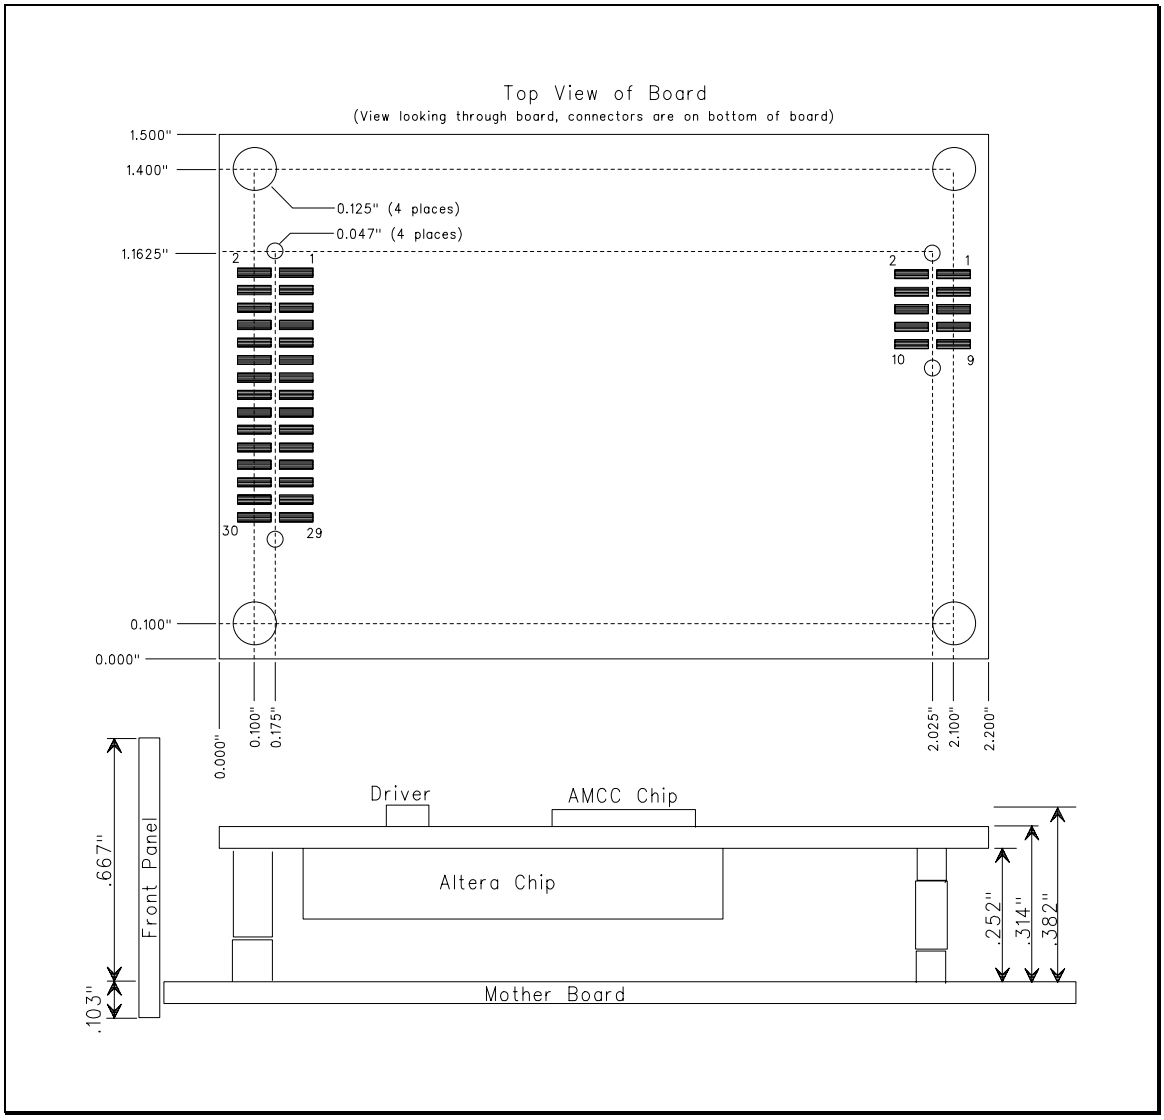

Fig. I-8. Mechanical dimensions.

Appendix A.

SERIAL TRANSMITTER DAUGHTER CARD SCHEMATIC DIAGRAM.

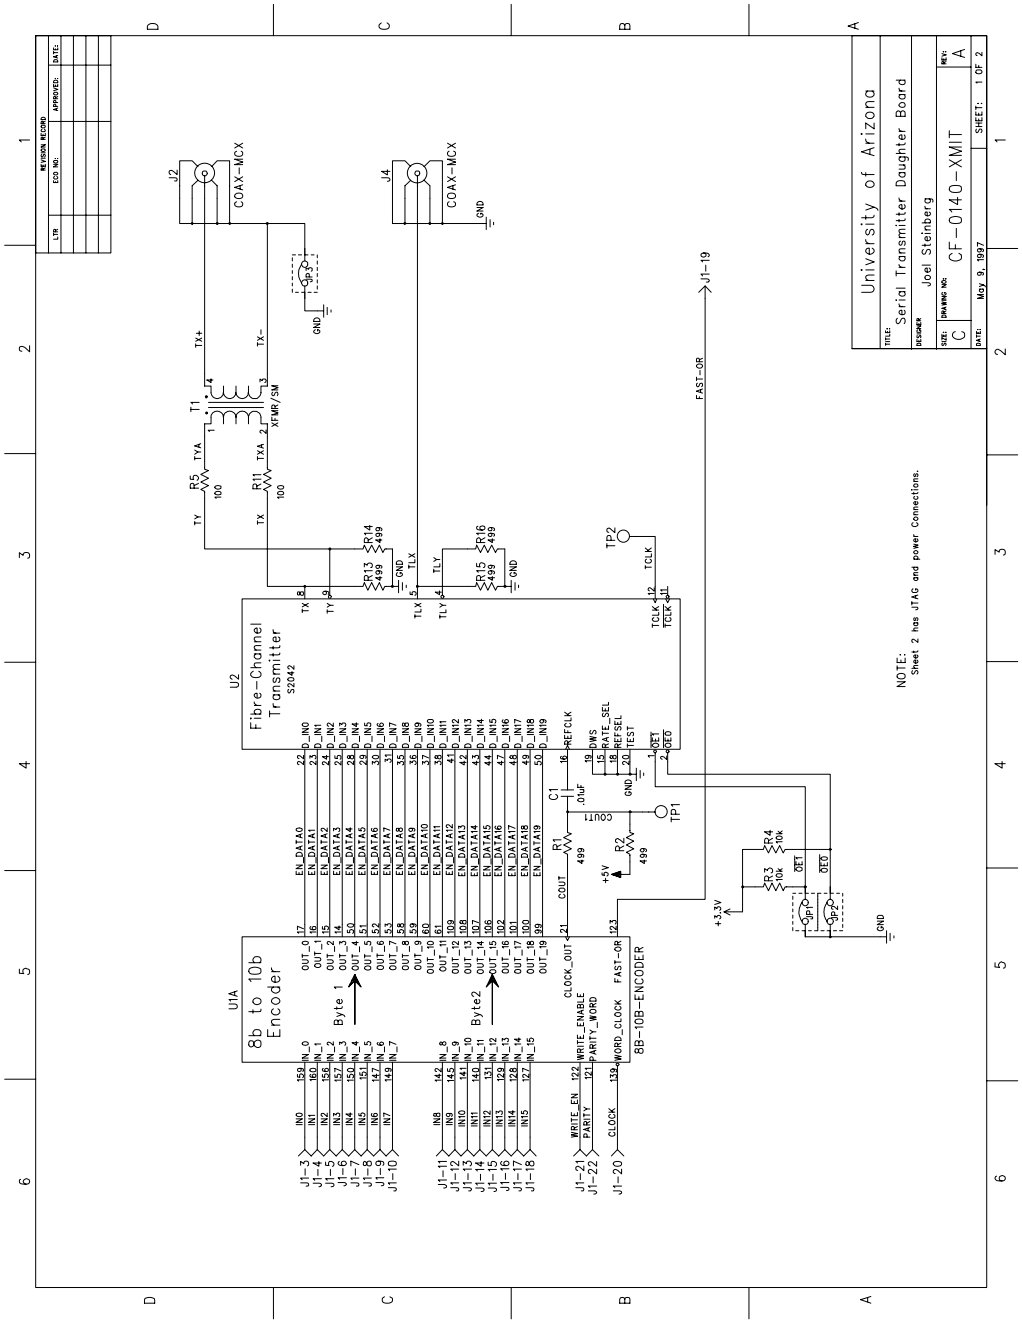

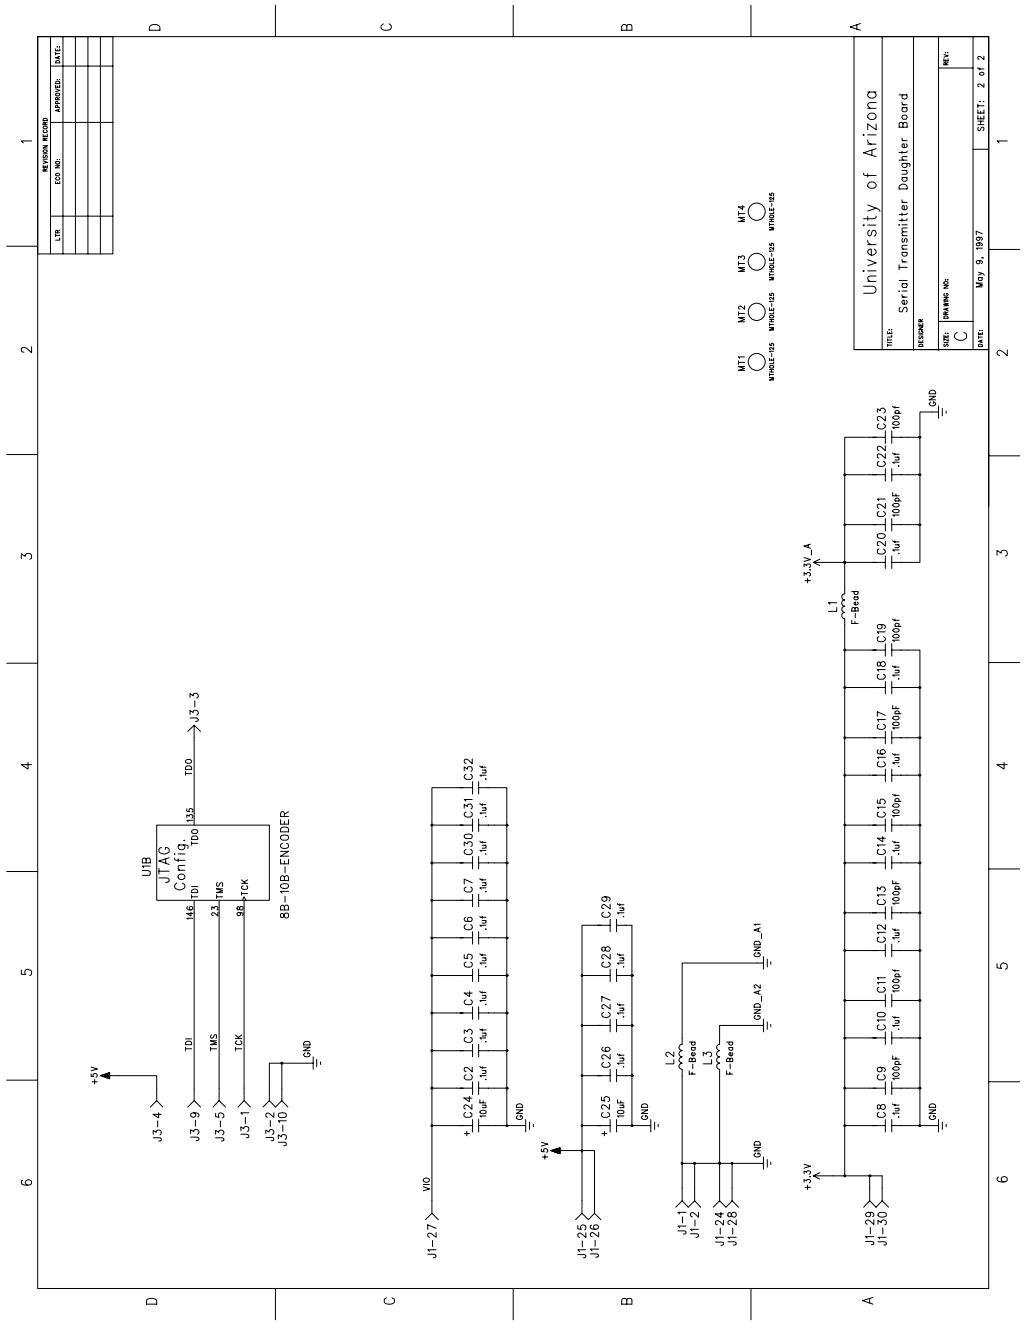
University of Arizona
Serial Transmitter Daughter Board
8B-10B-ENCODER
JTAG Config.
U1B
J3-4
J3-9
J3-5
J3-1
J3-2
J3-10
J3-3
TDI
TMS
TCK
TDO
+5V
GND
VIO
+3.3V
+3.3V_A
GND_A1
GND_A2
L1
L2
L3
F-Bead
J1-27
J1-25
J1-26
J1-1
J1-2
J1-24
J1-28
J1-29
J1-30
MT1
MT2
MT3
MT4
MTHOLE-125
REVISION RECORD
LTR
ECO NO:
APPROVED:
DATE:
SIZE:
C
DRAWING NO:
REV:
DATE:
May 9, 1997
SHEET: 2 of 2